

**Project Title:** Information Structure-Use Knowledge-Based Expert Systems  
**Principal Investigator:** Bruce Buchanan & Ted Shortliffe  
**Award Amount:** \$333,710  
**Period Covered:** 3/15/84 - 2/28/87

**Agency:** National Science Foundation  
**Project Title:** Interpretation of NMR Data for Proteins Using AI Methods  
**Principal Investigator:** Bruce Buchanan & Oleg Jardetzky  
**Award Amount:** \$100,000  
**Period Covered:** 11/1/84 - 10/31/86

**Agency:** National Science Foundation; MCS-8303142  
**Project Title:** The Mechanization of Formal Reasoning (Computer Research)  
**Principal Investigator:** Edward A. Feigenbaum  
**Award Amount:** \$286,144  
**Period Covered:** 7/15/83 - 6/30/86

**Agency:** Office of Naval Research  
**Project Title:** Computer-Based Tutors for Explaining and Managing the Process of Diagnostic Reasoning  
**Principal Investigator:** Bruce G. Buchanan  
**Award Amount:** \$513,662  
**Period Covered:** 3/15/85 - 3/14/88

**Agency:** Office of Naval Research [Pending]  
**Project Title:** Expert Control of Problem-Solving Search  
**Principal Investigator:** Bruce G. Buchanan  
**Award Amount:** \$725,899  
**Period Covered:** 8/1/85 - 7/31/88

**Agency:** Rockwell International [Pending]  
**Project Title:** Intelligent Real-Time Process Control  
**Principal Investigators:** Edward A. Feigenbaum  
**Amount:** \$643,059  
**Period Covered:** 1/1/87-12/31/89

**Agency:** Boeing Computer Services Company [Pending]  
**Project Title:** Research on Blackboard Systems and Intelligent Resource Management Systems  
**Principal Investigators:** E.A. Feigenbaum and B.G. Buchanan  
**Amount Requested:** \$569,604  
**Period Covered:** 4/1/86-3/31/88

**Agency:** Office of Naval Research [Pending]  
**Project Title:** Modeling Expert Control Knowledge  
**Principal Investigator:** Bruce G. Buchanan  
**Amount Requested:** \$96,909  
**Period Covered:** 11/1/85 - 10/31/86  
**Percentage Effort Committed to Project:** 3%, no salary support

**Agency:** Defense Advanced Research Projects Agency [Pending]  
**Project Title:** Research on Intelligent Assistants for Protein Chemists  
**Principal Investigator:** Bruce G. Buchanan, Oleg Jardetzky, Hans Andersen, Harden McConnell  
**Amount Requested:** \$7,962,678

**Period Covered:** 10/1/86-9/30/91

**Agency:** National Science Foundation [Pending]

**Project Title:** Heuristic Refinement for Deriving Solution Structures of Proteins

**Principal Investigator:** Bruce G. Buchanan and Oleg Jardetzky

**Amount Requested:** \$589,764

**Period Covered:** 11/1/86-10/31/89

#### **Interactions with the SUMEX-AIM Resource**

Our interactions with the SUMEX-AIM resource involve the facilities -- both hardware and software -- and the staff -- both technical and administrative. Taken together as a whole resource, they constitute an essential part of the research structure for the KSL. Many of the grants and contracts from other agencies have been awarded partly because of the cost-effectiveness of AI research in the KSL due to the fact that much of our computing needs could be more than adequately met by the SUMEX-AIM resource.

We rely on the central SUMEX facility as a focal point for all the research within the KSL, not only for much of our computing, but for communications and links to our many collaborators as well. As a common communications medium alone, it has significantly enhanced the nature of our work and the reach of our collaborations. The existence of the central time-shared facility has allowed us to explore new ideas at very small incremental cost.

As SUMEX and the KSL acquire a diversity of hardware, including LISP workstations machines and smaller personal computers, we rely more and more heavily on the SUMEX staff for integration of these new resources into the local network system. The staff has been extremely helpful and effective in dealing with the myriad of complex technical issues and leading us competently into this world of decentralized, diversified computing. At the same time, the staff has provided a stable, efficient central time-shared machine running software that has been developed at many sites over many years. Without the dedication of the SUMEX staff, the KSL would not be at the forefront of AI research.

### III.A.3.5. Training Activities

The SUMEX resource exists to facilitate biomedical artificial intelligence applications from program development through testing in the target research communities. This user orientation on the part of the facility and staff has been a unique feature of our resource and is responsible in large part for our success in community building. The resource staff has spent significant effort in assisting users gain access to the system and use it effectively. We have also spent substantial effort to develop, maintain, and facilitate access to documentation and interactive help facilities. The HELP and Bulletin Board subsystems have been important in this effort to help users get familiar with the computing environment.

On another front, we have regularly accepted a number of scientific visitors for periods of several months to a year, to work with us to learn the techniques of expert system definition and building and to collaborate with us on specific projects. Our ability to accommodate such visitors is severely limited by space, computing, and manpower resources to support such visitors within the demands of our on-going research.

Finally, the training of graduate students is an essential part of the research and educational activities of the KSL. Currently, 41 students are working with our projects centered in Computer Science and another 20 students are working with the Medical Computer Science program in Medicine. Of the 41 working in Computer Science, 25 are working toward Ph.D. degrees, and 16 are working toward M.S. degrees. A number of students are pursuing interdisciplinary programs and come from the Departments of Engineering, Mathematics, Education, and Medicine.

Based on the SUMEX-AIM community environment, we have initiated two unique and special academic degree programs at Stanford, the Medical Information Science program and the Masters of Science in AI, to increase the number of students we produce for research and industry, who are knowledgeable about knowledge-based system techniques.

The *Medical Information Sciences (MIS)* program is one of the most obvious signs of the local academic impact of the SUMEX-AIM resource. The MIS program received recent University approval (in October 1982) as an innovative training program that offers MS and PhD degrees to individuals with a career commitment to applying computers and decision sciences in the field of medicine. The MIS training program is based in the School of Medicine, directed by Dr. Shortliffe, co-directed by Dr. Fagan, and overseen by a group of nine University faculty that includes several faculty from the Knowledge Systems Laboratory (Profs. Shortliffe, Feigenbaum, Buchanan, and Genesereth). It was Stanford's active ongoing research in medical computer science, plus a world-wide reputation for the excellence and rigor of those research efforts, that persuaded the University that the field warranted a new academic degree program in the area. A group of faculty from the medical school and the computer science department argued that research in medical computing has historically been constrained by a lack of talented individuals who have a solid footing in both the medical and computer science fields. The specialized curriculum offered by the new program is intended to overcome the limitations of previous training options. It focuses on the development of a new generation of researchers with a commitment to developing new knowledge about optimal methods for developing practical computer-based solutions to biomedical needs.

The program accepted its first class of four trainees in the summer of 1983 and a second class of five entered last summer. A third group of seven students has just been selected to begin during 1985. The proposed steady state size for the program (which should be reached in 1986) is 20-22 trainees. Applicants to the program in our first two years have come from a number of backgrounds (including seven MD's and five medical students). We do not wish to provide too narrow a definition of what kinds of

prior training are pertinent because of the interdisciplinary nature of the field. The program has accordingly encouraged applications from any of the following:

- medical students who wish to combine MD training with formal degree work and research experience in MIS;
- physicians who wish to obtain formal MIS training after their MD or their residency, perhaps in conjunction with a clinical fellowship at Stanford Medical Center;
- recent BA or BS graduates who have decided on a career applying computer science in the medical world;
- current Stanford undergraduates who wish to extend their Stanford training an extra year in order to obtain a "co-terminus" MS in the MIS program;
- recent PhD graduates who wish post-doctoral training, perhaps with the formal MS credential, to complement their primary field of training.

In addition, a special one-year MS program is available for established academic medical researchers who may wish to augment their computing and statistical skills during a sabbatical break.

With the exception of this latter group, all students spend a minimum of two years at Stanford (four years for PhD students) and are expected to undertake significant research projects for either degree. Research opportunities abound, however, and they of course include the several Stanford AIM projects as well as research in psychological and formal statistical approaches to medical decision making, applied instrumentation, large medical databases, and a variety of other applications projects at the medical center and on the main campus. Several students are already contributing in major ways to the AIM projects and core research described in this application.

Early evidence suggests that the program already has an excellent reputation due to:

- high quality students, many of whom are beginning to publish their work in conference proceedings and refereed journals;
- a rigorous curriculum that includes newly-developed course offerings that are available to the University's medical students, undergraduates, and computer science students as well as to the program's trainees;
- excellent computing facilities combined with ample and diverse opportunities for medical computer science and medical decision science research;
- the program's great potential for a beneficial impact upon health care delivery in the highly technologic but cost-sensitive era that lies ahead.

The program has been successful in raising financial and equipment support (almost \$1M in hardware gifts from Hewlett Packard, Xerox, and Texas Instruments; over \$200K in cash donations from corporations and foundations; and an NIH post-doctoral training grant from the National Library of Medicine).

The *Master of Science in Computer Science: Artificial Intelligence (MS:AI)* program is a terminal professional degree offered for students who wish to develop a competence in the design of substantial knowledge-based AI applications but who do not intend to obtain a Ph.D. degree. The MS:AI program is administered by the Committee for Applied Artificial Intelligence, composed of faculty and research staff of the Computer Science Department. Normally, students spend two years in the program with their

time divided equally between course work and research. In the first year, the emphasis is on acquiring fundamental concepts and tools through course work and project involvement. During the second year, students implement and document a substantial AI application project.

### III.A.3.6. Resource Operations and Usage

The following data give an overview of various aspects of SUMEX-AIM resource usage. There are 5 subsections containing data respectively for:

1. Overall resource loading data (page 57).
2. Relative system loading by community (page 58).
3. Individual project and community usage (page 61).
4. Network usage data (page 69).
5. System reliability data (page 69).

For the most part, the data used for these plots cover the entire span of the SUMEX-AIM project. This includes data from both the KI-TENEX system and the current DECsystem 2060. At the point where the SUMEX-AIM community switched over to the 2060 (February, 1983), you will notice sharp changes in most of the graphs. This is due to many reasons briefly mentioned here:

1. Even though the TENEX operating system used on the KI-10 was a forerunner of the current Tops20 operating system, the Tops20 system is still different from TENEX in many ways. Tops20 uses a radically different job scheduling mechanism, different methods for computing monitor statistics, different I/O routines, etc. In general, it can not be assumed that statistics measured on the TENEX system correlate one to one with similar statistics under Tops20.
2. The KL-10 processor on the 2060 is a faster processor than the KI-10 processor used previously. Hence, a job running on the KL-10 will use less CPU time than the same job running on the KI-10. This aspect is further complicated by the fact that the SUMEX KI-10 system was a dual processor system.
3. The SUMEX-AIM Community was changing during the time of the transfer to the 2060. The usage of the GENET community on SUMEX had just been phased out. This part of the community accounted for much of the CPU time used by the AIM community. Since the purchase of the 2060 was partially funded by the Heuristic Programming Project (HPP), an additional number of HPP Core Research Projects started using the 2060, increasing the Stanford communities usage of the machine.
4. Finally, the move to the 2060 occurred during a pivotal time in the community when more and more projects were either moving to their own local timesharing machines, or onto specialized Lisp workstations. It also was the time for the closure of many long time SUMEX-AIM projects, like DENDRAL and PUFF/VM.

Any conclusions reached by comparing the data before and after February, 1983 should be done with caution. The data is included in this year's annual report mostly for casual comparison.

*Overall Resource Loading Data*

The following plot displays total CPU time delivered per month. This data includes usage of the KI-TENEX system and the current DECsystem 2060.




Figure 7: Total CPU Time Consumed by Month

*Relative System Loading by Community*

The SUMEX resource is divided, for administrative purposes, into three major communities: user projects based at the Stanford Medical School (*Stanford Projects*), user projects based outside of Stanford (*National AIM Projects*), and common system development efforts (*System Staff*). As defined in the resource management plan approved by the BRP at the start of the project, the available system CPU capacity and file space resources are divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems, documentation, etc.

The monthly usage of CPU resources and terminal connect time for each of these three communities relative to their respective aliquots is shown in the plots in Figure 8 and Figure 9. As mentioned on page 56, these plots include both KI-10 and 2060 usage data.






Figure 8: Monthly CPU Usage by Community




Figure 9: Monthly Terminal Connect Time by Community

*Individual Project and Community Usage*

The following histogram and table show cumulative resource usage by collaborative project and community during the past grant year. The histogram displays the project distribution of the total CPU time consumed between May 1, 1985 and April 30, 1986, on the SUMEX-AIM DECsystem2060 system.

In the table following, entries include a text summary of the funding sources (outside of SUMEX-supplied computing resources) for currently active projects, total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May, 1985 and May, 1986.




Figure 10: Cumulative CPU Usage Histogram by Project and Community

## Resource Use by Individual Project - 5/85 through 4/86

<i>National AIM Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) CADUCEUS "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. Randolph A. Miller, M.D. University of Pittsburgh NIH 5 R24 RR-01101-08 7/80-6/86 \$1,658,347 7/84-9/85 \$354,211 10/85-3/86 \$50,690 NIH 5 R01 LM03710-05 7/80-3/86 \$853,200 7/84-9/85 \$210,091 10/85-3/86 \$35,316 NIH New Invest 5 R23 LM03889-03 Gordon E. Banks, M.D. 4/82-3/85 \$107,675 4/84-3/85 \$35,975 NIH 1 KO4 LM00084-01 Randolph A. Miller, M.D. 9/85-9/86 \$55,296 10/86-9/90 to be determined annually	39.32	1274.18	1326
2) CLIPR Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado NIMH 5 R01 MH-15872-14-16 (Kintsch) 7/84-6/87 \$145,500 7/83-6/84 \$56,501 NSF (Kintsch) 8/83-7/86 \$200,000(*) IBM (Polson) David Kieras University of Arizona 1/85-12/86 \$500,000 1/86-12/86 \$250,000	0.75	122.33	144

3) SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota Control Data Corp. (Johnson) 1983-85 \$90,000 1986-88 \$95,000 Microelect. and Info. Ctr. Univ. of MN (Johnson, Thompson, Slagle, Wechsler, Yonas) 1984-1985 \$500,000 1985-1986 \$300,000 NIH LM-00160 (Johnson, Connelly) 1984-1989 \$712,573 McKnight Foundation (Johnson, Bailey) 1984-1985 \$13,000 Dwan Family Fund, Univ. of MN Medical School (Johnson) 1985 \$6,000	2.06	301.35	650
4) MENTOR Project "Medical Evaluation of Therapeutic Orders" Stuart M. Speedie, Ph.D. University of Maryland Terrence F. Blaschke, M.D. Stanford University National Center for Health Services Research 1 R18 HS 05263 1/85-12/88 \$485,134 1/86-12/86 \$182,820	27.28	1816.13	650
5) *** [Rutgers-AIM] *** Rutgers Research Resource Artificial Intelligence in Medicine Casimir Kulikowski, Ph.D. Sholom Weiss, Ph.D. Rutgers U., New Brunswick NIH RR-02230-02 (Kulikowski, Weiss) 12/83-11/87 \$3,198,075 12/84-11/85 \$613,897	0.03	3.79	0
6) AIM Pilot Projects	51.92	4754.44	394
7) AIM Administration	0.36	43.92	110
8) AIM Users	27.80	4007.09	2423
	-----	-----	-----
Community Totals	149.52	11838.18	5555

<i>Stanford Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) GUIDON-NEOMYCIN Project Bruce G. Buchanan, Ph.D. William J. Clancey, Ph.D. Dept. Computer Science ONR/ARI N00014-79-C-0305 3/79-3/86 \$510,311(*) Josiah Macy, Jr. Foundation 3/85-3/88 \$503,415	59.02	10173.89	1432
2) MOLGEN Project "Applications of Artificial Intelligence to Molecular Biology: Research in Theory Formation, Testing and Modification" Edward A. Feigenbaum, Ph.D. Peter Friedland, Ph.D. Charles Yanofsky, Ph.D. Depts. Computer Science/ Biology NSF MCS-8310236 (Feigenbaum, Yanofsky) 11/83-10/85 \$270,836(*) 11/85-10/86 \$135,000(*)	57.66	7121.50	1432
3) ONCOCIN Project "Knowledge Engineering for Med. Consultation" Edward H. Shortliffe, M.D., Ph.D. Dept. Medicine NIH RR-01613 7/83-6/86 \$624,455 7/85-6/86 \$220,302 NIH LM-04136 8/83-7/86 \$211,851 8/85-7/86 \$74,150 H.J. Kaiser Family Fdn. 7/83-6/88 \$200,000 7/85-6/86 \$50,000 NSF IST83-12148 Bruce G. Buchanan (Shortliffe) 3/84-2/87 \$330,000(*) 3/85-2/86 \$52,679(*) NIH 1 T32 LM07033 7/84-6/89 \$903,718 7/85-6/86 \$215,850 NIH 1 R23 LM04316 2/85-1/88 \$107,441 2/85-1/86 \$37,500 NIH 1 R01 LM0442001 9/85-8/88 \$314,707 9/85-8/86 \$95,205	143.47	242629.39	2619

IST 83-12148			
3/84-2/87 \$330,000(*)			
3/85-2/86 \$52,679			
4) PROTEAN PROJECT	223.52	8577.51	4368
Oleg Jardetzky			
School of Medicine			
Bruce Buchanan			
Computer Science Department			
NSF PCM-84-02348			
11/84-10/86 \$100,000(*)			
11/85-10/86 \$50,000(*)			
5) RADIX Project	22.47	2918.85	343
"Deriving Medical Knowledge from			
Time Oriented Clinical Databases"			
Robert L. Blum, M.D.			
Gio C.M. Wiederhold, Ph.D.			
Depts. Computer Science/			
Medicine			
NSF IST-83-17858 (Blum)			
3/84-3/86 \$89,597(*)			
NIH LM-04334 (Wiederhold)			
5/84-11/86 \$291,192			
6) Stanford Pilot Projects	54.25	3468.37	1964
7) Core AI Research	65.68	8610.18	1263
8) Stanford Associates	4.19	440.39	1100
9) Medical Information Sciences	62.19	8227.63	736
	-----	-----	-----
Community Totals	689.19	74167.62	16777



<i>KSL-AI Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
For funding details please see page 49			
1) Advanced Architectures	64.78	24732.17	1663
2) FOL	142.06	5528.15	1439
3) Intelligent Agent	119.41	13320.77	2840
4) Pixie	35.31	4021.73	875
5) KSL Management	80.02	15484.89	2837
6) DART	11.19	2906.89	811
7) MRS	96.96	8503.16	1456
8) Helix	19.18	7287.03805	
	-----	-----	-----
Community totals	504.22	77586.54	14210
<i>SUMEX Staff</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	280.73	30889.71	9810
2) System Associates	17.83	1235.53	1529
	-----	-----	-----
Community Totals	303.27	32686.88	12129
<i>System Operations</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	752.65	83590.22	9902
	=====	=====	=====
Resource Totals	2397.04	279869.50	58573

(\*) Award includes indirect costs.

*System Reliability*

System reliability for the DECsystem 2060 has significantly improved in this past period. We have had very few periods of particular hardware or software problems. The data below covers the period of May 1, 1985 to April 30, 1986. The actual downtime was rounded to the nearest hour.

Table 1 : System Downtime Hours per Month - May 1985 through April 1986

8	1	11	9	26	12	1	22	1	7	18	22
May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr

Table 2 : System Downtime Hours per Month - May 1985 through April 1986

Reporting period	:	365 days, 0 hours, 12 minutes, and 49 seconds
Total Up Time	:	358 days, 6 hours, 37 minutes, and 16 seconds
PM Downtime	:	1 days, 8 hours, 49 minutes, and 50 seconds
Actual Downtime	:	5 days, 8 hours, 45 minutes, and 43 seconds
Total Downtime	:	6 days, 17 hours, 35 minutes, and 33 seconds
Mtbf	:	2 days, 14 hours, 45 minutes, and 49 seconds
Uptime Percentage	:	98.53

*Network Usage Statistics*

The plots in Figure 11 and Figure 12 show the monthly network terminal connect time for the TYMNET and the INTERNET usage. The INTERNET is a broader term for what was previously referred to as Arpanet usage. Since many vendors now support the INTERNET protocols (IP/TCP) in addition to the Arpanet, which converted to IP/TCP in January of 1983, it is no longer possible to distinguish between Arpanet usage and Internet usage on our 2060 system.




Figure 11: TYMNET Terminal Connect Time




Figure 12: ARPANET Terminal Connect Time

### III.A.4. Future Plans

Our plans for the next grant year (year 14) are based on the Council-approved plans for our 3-year renewal to begin in August, 1986. In the next period, we will merge into SUMEX-AIM two on-going BRTP-supported efforts: the on-going SUMEX-AIM work on biomedical AI applications and the ONCOCIN work on studies in the dissemination of consultation systems. The directions and background for much of this work derive from our research to date.

The goals of the SUMEX-AIM resource are long term in supporting basic research in artificial intelligence, applying these techniques to a broad range of biomedical problems, developing the methodologies for disseminating AI systems into the biomedical community, experimenting with communication technologies to promote scientific interchange, and developing better tools and facilities to carry on this research.

#### Toward a More Distributed Resource

In the early 1970's, the initial model for SUMEX-AIM as a centralized resource was based on the high cost of powerful computing facilities and the infeasibility of being able to duplicate them readily. As planned, this central role has already evolved significantly and continues to evolve with the introduction of more compact and inexpensive computing technology now available at many more research sites. At the same time, the number of active groups working on biomedical AI problems has grown and the established ones have increased in size. This has led to a growth in the demand for computing resources far beyond what SUMEX-AIM could reasonably and effectively provide on a national scale. We have actively supported efforts by the more mature AIM projects to develop or adapt additional computing facilities tailored to their particular needs and designed to free the main SUMEX resource for new, developing applications projects. To date, over 10 of the national projects have moved some or all of their work to local sites and several have begun resource communities of their own (see page 87). Thus, as more remotely available resources have become established, the balance of the use of the SUMEX-AIM resource has shifted toward supporting start-up pilot projects and the growing AI research community at Stanford.

#### Summary of Specific Objectives

Our future goals for the central SUMEX-AIM resource are then guided by:

- The increasingly decentralized character of the resource and community and the need to find ways to maintain the scientific communication and sharing that has characterized SUMEX-AIM work
- The continuing exploration of important new areas of biomedical research in which AI techniques can be effectively applied
- The need for a strong basic research effort to investigate methodologies to attack the many problems still beyond our current AI systems and to develop improved tools to build more complex and effective expert systems
- The growing impact of biomedical AI and the need to find and evaluate ways for effectively disseminating biomedical AI technology into real-world settings.
- The need for computing environments for our research and dissemination work that anticipate the needs of AI applications systems over the next 5-10

years, based on the rapidly changing computing hardware and software technology base

SUMEX-AIM will retain its role as a national experimental laboratory for biomedical AI research with a double thrust -- on the one hand, pursuing the basic research for, experimentation on, and trial dissemination of interesting applications and on the other, anticipating and developing the model computing and community environment in which this work can take place. We will nurture existing and new projects and serve as a communications cross-roads for the now diverse AIM community. We will provide the computing resources and some manpower support for long-term basic AI research activities that promise to illuminate issues relevant to future selected collaborative application areas in biology and medicine. For example, our detailed plans have many threads between our basic research in general patient treatment protocol acquisition, representation, and decision-making tools and our collaborative applications in cancer chemotherapy or hypertension trials. Or between our basic research in blackboard problem-solving frameworks and system architectures and our collaborative application in NMR protein conformation determination. Other basic research areas have even longer term goals for problems we hope to be able to address in the future. Underlying all this work will be the development of the Lisp workstation system and network environment that will facilitate these research results and that we feel will become the routine computing environment of the next decade.

In all of this, SUMEX will be both a working laboratory for selected projects within our computing and manpower capacity limits and a source and repository for software and ideas for a broader remote community. We will become an increasingly distributed community resource with heterogeneous computing facilities tethered to each other through various communications media. Many of these machines will be located physically near the projects or biomedical scientists using them. We retain our sincere commitment to our national community of projects. But, inevitably their needs will be met more and more by local facilities and our plans as a resource for the next term place greater emphasis than in the past on supporting the growing Stanford community of AIM collaborations and projects and on developing and integrating model systems at Stanford that can be emulated elsewhere for AIM community needs.

Even with more distributed computing resources, the central resource will continue to play an important role for the next term as a communication crossroads and as a focus for our active dissemination efforts. A key challenge will be to maintain the scientific community ties that grew naturally out of the previous co-location within a central facility.

The following outlines the specific objectives of the SUMEX-AIM resource during the follow-on five year period. Note that these objectives cover only the resource nucleus; objectives for individual collaborating projects are discussed in their respective reports in Section IV. Specific aims are broken into five categories: 1) Core Research and Development, 2) Collaborative Research, 3) Service and Resource Operations, 4) Training and Education, and 5) Dissemination.

#### *1) Core Research and Development*

SUMEX funding and computational support for core research is complementary to similar funding from other agencies (see page 49) and contributes to the long-standing interdisciplinary effort at Stanford in basic AI research and expert system design. We expect this work to provide the underpinnings for increasingly effective consultative programs in medicine and for more practical adaptations of this work within emerging microelectronic technologies. Specific aims include:

- Basic research on AI techniques applicable to biomedical problems. Over

the next term we will emphasize work on blackboard problem-solving frameworks and architectures, knowledge acquisition or learning, constraint satisfaction, and qualitative simulation.

- Investigate methodologies for disseminating application systems such as clinical decision-making advisors into user groups. This will include generalized systems for acquiring, representing and reasoning about complex treatment protocols such as are used in cancer chemotherapy and which might be used for clinical trials.
- Support community efforts to organize and generalize AI tools and architectures that have been developed in the context of individual application projects. This will include retrospective evaluations of systems like the AGE blackboard experiment and work on new systems such as BB1, MRS, SOAR, EONCOCIN, EOPAL, Meta-ONYX, and architectures for concurrent symbolic computing. The objective is to evolve a body of software tools that can be used to more efficaciously build future knowledge-based systems and explore other biomedical AI applications.
- Develop more effective workstation systems to serve as the basis for research, biomedical application development, and dissemination. We seek to coordinate basic research, application work, and system development so that the AI software we develop for the next 5-10 years will be appropriate to the hardware and system software environments we expect to be practical by then. Our purchases of new hardware will be limited to experimentation with state-of-the-art workstations as they become available for our system developments.

## *2) Collaborative Research*

- Encourage the exploration of new applications of AI to biomedical research and improve mechanisms for inter- and intra-group collaborations and communications. While AI is our defining theme, we may consider exceptional applications justified by some other unique feature of SUMEX-AIM essential for important biomedical research. We will continue to exploit community expertise and sharing in software development.
- Minimize administrative barriers to the community-oriented goals of SUMEX-AIM and direct our resources toward purely scientific goals. We will retain the current user funding arrangements for projects working on SUMEX facilities. User projects will fund their own manpower and local needs; actively contribute their special expertise to the SUMEX-AIM community; and receive an allocation of computing resources under the control of the AIM management committees. There will be no "fee for service" charges for community members.
- Provide effective and geographically accessible communication facilities to the SUMEX-AIM community for remote collaborations, communications among distributed computing nodes, and experimental testing of AI programs. We will retain the current ARPANET and TYMNET connections for at least the near term and will actively explore other advantageous connections to new communications networks and to dedicated links.

### *3) Service and Resource Operations*

Our applications, core research, and system development will be directed toward realizing and exploiting the computing environment that will be routinely available in the late 1980's and early 1990's, based on compact, decentralized, high-performance personal workstations that take advantage of the intelligent computing environments beginning to emerge from today's Lisp workstations. Consistent with these plans, we will immediately discontinue BRTP subsidy for the DEC 2020 demonstration machine and for the shared VAX 11/780 time-sharing system. Also we will gradually and responsibly phase out BRTP support for the DEC 2060 mainframe system that has been our chief shared resource and link to the past.

SUMEX-AIM does not have the computing or manpower capacity to provide routine service to the large community of mature projects that has developed over the years. Rather, their computing needs are better met by the appropriate development of their own computing resources when justified. Thus, SUMEX-AIM has the primary focus of assisting new start-up or pilot projects in biomedical AI applications in addition to its core research in the setting of a sizable number of collaborative projects. We do offer continuing support for projects through the lengthy process of obtaining funding to establish their own computing base.

### *Training and Education*

- Provide documentation and assistance to interface users to resource facilities and systems.
- Exploit particular areas of expertise within the community for assisting in the development of pilot efforts in new application areas.
- Accept visitors in Stanford research groups within limits of manpower, space, and computing resources.
- Support the Medical Information Science and MS/AI student programs at Stanford to increase the number of research personnel available to work on biomedical AI applications.
- Support workshop activities including collaboration with other community groups on the AIM community workshop and with individual projects for more specialized workshops covering specific research, application, or system dissemination topics.

### *5) Dissemination*

While collaborating projects are responsible for the development and dissemination of their own AI systems and results, the SUMEX resource will work to provide community-wide support for dissemination efforts in areas such as:

- Encourage and support the on-going export of software systems and tools within the AIM community and for commercial development.
- Assist in the production of video tapes and films depicting aspects of AIM community research.
- Promote the publication of books, review papers, and basic research articles on all aspects of SUMEX-AIM research.

### Resource Scope

The SUMEX-AIM resource has been from its inception a national experimental resource for biomedical AI with a scope that is carefully defined. Within its limited manpower and computational resources, its focus has been on experiments in new and varied biomedical applications of AI, assisting new research groups in biomedical AI get started, exploring ways to disseminate AI systems into biomedical user communities, supporting relevant basic AI research, and facilitating scientific communications and community sharing. The SUMEX-AIM user community comprises projects from many biological and medical disciplines, ranging from chemistry to molecular biology to clinical medicine to cognitive psychology, and represents collaborations between computer and biomedical scientists from many parts of Stanford University and other universities around the country. The development of this diverse community of projects has both justified the cost of and made effective use of SUMEX-AIM computational and communication facilities at Stanford and elsewhere in our resource community. In its resource role, SUMEX has intentionally limited its production computational capacity to meet the needs of national start-up projects and Stanford research groups, while encouraging self-sufficient community members to develop resources to meet their own computing needs. This has allowed us to provide a level of support for on-going projects and to concentrate most of our efforts on experiments with integrating emerging hardware and software technologies that will be the vehicles of future biomedical AI systems. The results of these experiments are widely disseminated and help other groups through example and direct export of software and ideas.



### **III.B. Highlights**

In this section we describe several research highlights from the past year's activities. These include notes on existing projects that have passed important milestones, new pilot projects that have shown progress in their initial stages, and some other special activities that reflect the impact and influence the SUMEX-AIM resource has had in the scientific and educational communities.